

For Reference

NOT TO BE TAKEN FROM THIS ROOM

Ex LIBRIS
UNIVERSITATIS
ALBERTAENSIS



THE UNIVERSITY OF ALBERTA

CONVENTIONS AND CONVENTIONAL MEANING

by

RODERICK SCOTT WILTSHIRE

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES AND RESEARCH
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE
OF MASTER OF ARTS

DEPARTMENT OF PHILOSOPHY

EDMONTON, ALBERTA


SPRING, 1983

ABSTRACT

David Lewis has analyzed our common concept of convention, and has found conventions depend for their being so on the possibility of making certain complex psychological attributions to the persons involved in them. This has been understood by Jonathan Bennett to aid his project of giving an account of conventional meaning, including language, in terms which make only eliminable reference to the common, accepted or conventional meaning of an utterance. It has also been shown by Lewis that his analysis implies that utterances that are in accordance with communicative conventions are accurately described by H.P. Grice's analysis of non-natural meaning.

In this thesis, Lewis's development of his analysis is criticized, and counterexamples are adduced to demonstrate the roles of habit and practical reasoning in conventional behavior are misrepresented by that analysis. Alterations to Lewis's analysis which are intended to remedy these difficulties are introduced. Consideration of Jonathan Bennett's attempt to describe how behavioral evidence might be found for Lewis's claims about conventions finds this attempt to be unsuccessful. On the basis of this, it is argued that the criticisms of Lewis's analysis are corroborated, and a major tenet of Bennett's "Meaning Nominalism" is undermined. Lewis's application of his analysis to language is then discussed, and a counterexample to Lewis's claim that knowledge of alternative possible

conventions is essential to conventionality is taken to suggest that the sense in which language is conventional is not devoid of metaphor. Finally, Lewis's claim that conventional meaning is Gricean is shown to be unfounded, and a demonstration that conventional meaning is not Gricean is presented. Instances of the operation of the "Gricean Mechanism" in circumstances provided by normal language use are described, and some natural language phenomena which have been taken to demonstrate the operation of the Gricean Mechanism are shown to be falsely so thought.



Digitized by the Internet Archive
in 2023 with funding from
University of Alberta Library

<https://archive.org/details/Wiltshire1983>

Preface

I could call God "Charlie" and my dog Charlie a god if I chose, and if everyone else joined me in this no one would say or be thought to say anything silly or insulting to another's religion. I could drive on the left if I chose, and if everyone else did too no one would collide (at least, any more often than they do now). Driving on the right is a convention we have; language seems to be similar to driving on the right in significant ways. David Lewis has attempted to say exactly what it is about driving on the right and speaking as we do that makes them conventional.¹

Though the bulk of what follows takes the form of an attack on the adequacy of Lewis's analysis of convention, I take it to be essentially a development of that analysis in ways generally consistent with Lewis's intentions. The issues I want to emphasize in regards to conventions are things that Lewis is often at pains to point out. His analysis, because he accepts certain false analogies suggested by the way he develops it, does not adequately reflect important aspects of his own intuitions.

Lewis intends his analysis to subvert the idea that a convention must originate in an agreement. In the introduction to his Convention, he says that his analysis is a "theory along the lines of Hume's, in his discussion of the origin of justice and property. Convention turns out to be

¹David K. Lewis, Convention: A Philosophical Study (Cambridge, MA: Harvard University Press, 1969).

a general sense of common interest; which sense all the members of the society express to one another, and which induces them to regulate their conduct by certain rules. I observe that it will be in my interest [e.g.] to leave another in the possession of his goods, provided he will act in the same manner with regard to me. When this common sense of interest is mutually expressed and known to both, it produces a suitable resolution and behavior. And this may be properly enough be called a convention or agreement betwixt us, though without the interposition of a promise; since the actions of each of us have a reference to those of the other, and are performed upon the supposition that something is to be performed on the other part." (interpolation Lewis's)²

All of this is well and good. Conventions depend upon a certain kind of community of interest and action among people.

My concern can be articulated by pointing out a divergence between Hume's and Lewis's practice. Both describe the conditions for conventionality in a manner appropriate to the description of planned action. This serves to mark in an intelligible manner the ecological relationships essential to conventionality. When Hume describes how it is that we conform to convention he does not in general continue in this manner. Our conventional action--in Hume's discussion, doing what is just, and respecting other's property--becomes a matter of following a general rule, motivated not by an understanding of what is in our interests in the particular case, but by a developed capacity for appreciating and motivation to do what is in the common interest. For Hume, it is specifically not a

²Ibid., p.3f. Hume quoted from A Treatise of Human Nature, III.ii.2.

process of practical reasoning that induces us to do what is just. Lewis does not follow Hume in this, partly perhaps because the independence of convention from action demanded by practical evaluation of gains and losses is not so obvious in the case of driving on the right as it is in the case of just action. Lewis describes our action in accordance with convention as if it were motivated by reasoning analogous to that used in the "ecological" model of conventional action. This is the source of what concerns me about Lewis's analysis. Though he is at pains to deny there need be a distinction between "habitual" action and "rational" action (a good Humean point, this), and to deny his analysis implies the reasoning it apparently entails actually goes on, his practice gives aid and comfort to two doctrines current in the Philosophy of Language which are inconsistent with these denials.

One of these doctrines is the application of H.P. Grice's analysis of meaning to conventional meaning, which implies that one means what one does, and by implication one is understood, in virtue of a certain complicated set of intentions and inferences. The other, parasitic on the first, is Jonathan Bennett's meaning-nominalism, the claim that conventional meaning can only be properly understood in terms of the individual case of meaningful utterance, and that this in turn is best explicated in terms of Grice's analysis.

In Section One below, I address problems in Lewis's development of his analysis of convention that lead to the inadequacies that are diagnosed and remedied in Section Two. In Section Three, consideration of Bennett's attempt to justify an observationally-based version of aspects of Lewis's analysis which I claim to be faulty corroborate the argument of Section Two and serve to refute the major contention of meaning-nominalism. In Section Four I attempt to articulate a difference between two classes of conventions, and describe how Lewis's analysis misleads with respect to the members of one of these classes, one being language. In Section Five Grice's analysis of Non-Natural Meaning is shown not to apply to instances of meaning in virtue of communicative conventions, and the conditions under which normal language use results in situations to which Grice's analysis applies are described. In Section Six, some phenomena of normal language use which have been taken to suggest the appropriateness of Grice's analysis are shown to be falsely so thought.

Table of Contents

Chapter	Page
I. Co-ordination Situations and Their Solutions	1
II. Lewis's Problem-Solving Model of Conventional Activity	14
III. Behavioral Evidence for Higher-Order Reasons Underlying Conventions	30
IV. The Conventionality of Language	39
V. Conventional Meaning and Speaker's Intentions	50
VI. The Gricean Mechanism in Conventional Language Use	58
VII. Concluding Remarks	65
VIII. SELECT BIBLIOGRAPHY	67

I. Co-ordination Situations and Their Solutions

David Lewis develops the argument of his Convention (C) through consideration of a certain class of situations called co-ordination problems. As Lewis attempts to give a precise articulation of what a co-ordination problem is in game-theoretic terms, and as the game-theoretical description is not, in the end, essential to the analysis of the concept of convention that results, I will instead speak of co-ordination situations (except when considerations of euphony override) and depend on the intuitive similarity of the members of the class in my exposition of Lewis's view.

A co-ordination situation is one that is similar to the following. Two people have been talking on the telephone and have been disconnected. If they wish to continue their conversation, there are four possible combinations of the things they might do. They might both call the other back, and each find the line busy; they might each wait for the other to call back and wait indefinitely; one might call back while the second waits, permitting them to resume their conversation; or, the second might call back while the first waits, again enabling them to continue their conversation. What is to be noticed about such a situation is that each person's success depends on what the other does, and if and only if the persons involved can co-ordinate their actions can they both achieve their goals. What acts as an obstacle to the simple solution of such a problem is the fact that there is more than one way for the persons involved to "get

together", and the action required of each person is different for the different possible preferred combinations of actions. If, in our example, one of our callers was using a defective telephone and could not call out, and both persons knew this, there would be no problem. If the conversation is to continue, the situation decides which person must call and which must wait. Other examples of co-ordination situations are common. If two (or more) of us wish to meet, it is often the case that precisely where or when we meet is not a significant consideration, so long as we all attempt to meet the other(s) at the place and time the others do. If we are attempting to find something in a hurry, it does not in general matter where of the places that need to be searched each of us looks, so long as we all look in different places and all the places that need to be searched are searched by someone. If we are passing each other in the hall, it matters little if I pass you on the right or the left, so long as you do the same as I. If we are to run a three-legged race together, we can avoid falling all over each other by both starting with our outside legs or both with our inside legs.

In order to co-ordinate their actions, actors in a co-ordination situation must decide what to do on the basis of what they expect the other persons involved in the situation to do. As such, co-ordination must be achieved either by luck, or on the basis of the properties of the situation that give the person's involved a clue as to what

the others will do.

One way in which each person in a co-ordination situation might obtain strong reason to expect a course of action on the part of others is by agreeing with them as to what will be done. An explicit agreement gives them reason to expect that the others will do their part in the co-ordination, because co-ordination problems are of the kind we have specified; unilateral deviation from the agreement will result in events turning out badly for all persons involved.

The effects of agreement on the possibility of co-ordination points to a difference between the examples of co-ordination situations given above and one of the examples given by Lewis in the first section of Convention. If our livelihood depends on the success of our hunting, we may find ourselves in the following situation: We may choose to hunt deer together, and eat well, or we may each hunt alone, and hunt rabbit, and eat poorly. If some of us hunt deer, and others hunt rabbit, those hunting deer will starve, as everyone's participation is necessary for a successful deerhunt (C, p.7). In this example the second "co-ordination", that of each of us hunting rabbits, does not involve our "getting together" in the way that the other co-ordinations we have discussed do, for if someone decides to hunt rabbits, he need not prefer that everyone else does too.¹

¹The difference between cases similar to the deerhunt and other co-ordination problems is made clear in Section 5.3 of

The possibility of (somewhat) successful unilateral deviation from an agreement to hunt deer leaves a problem that does not exist in cases like our other examples, when they have been solved by agreement. In our other examples, the success of each person involved is linked to the success of the others involved in such a way that an agreement is a decisive reason to conform to what is agreed. In the deerhunt case, if my main concern is to avoid the worst that could happen (i.e. starvation), unless I am absolutely sure the others will all hunt deer, I will hunt rabbit, agreement or no. If I think others might feel the same way, and think that they might have concerns about whether I will show up for the deerhunt, I might be undecided as to what they might do. It may be that the slightest indecision on my part as to what the others will do will be reason enough for me to hunt rabbit. And this means, if I think there is the slightest likelihood of anyone else not hunting deer, I will hunt rabbit.² In this case, an agreement may exist, and I may still be faced with a problem to solve. The importance of this kind of case with respect to Lewis's analysis of convention will be discussed below. To mark the distinction between the deerhunt case and our other cases, I will say there is strong coincidence of interest in the other cases.

¹(cont'd)Edna Ullmann-Margalit's The Emergence of Norms (Oxford: Oxford University Press, 1977).

²Margaret Gilbert, in "Convention and Game Theory" (Synthese, Vol.46, 1981) argues that situations similar to this, wherein an agent can observe a strategy of avoiding the worst, raise significant problems for Lewis's attempted game-theoretic definition of co-ordination problems.

I take this to imply that we prefer a course of action only under the condition that the others take action co-ordinated with it.

In the event that communication between the persons involved in a co-ordination situation is impossible, the persons involved in the situation may be able to come to conclusions about what the others will do, and thereby decide what to do themselves, because one of the preferred combinations of actions is particularly salient. That is, one of the preferred combinations of actions may have some property that sets it apart from the others, and this property might be expected by all to be noticed by all, this common reference point may make it possible for co-ordination to be achieved. Here is an example of how this might work: If it is important that we meet, and there are three places that are candidates for our meeting, and if we have found two of them to be pleasant places to meet and the other to be noticeably less so, and if we are aware of the other's evaluations of the three meeting places, we may both reason thus: "If I go to either of the places that we would prefer to meet, I run the risk of missing you, because those two places do not distinguish themselves on the basis of what we both know about our problem. Since I know you are a smart person I know you know this too. Since I'm fairly sure you think I am a smart person, I'm fairly sure you know I know this too. The less pleasant place to meet is the only place that distinguishes itself from the others on the basis

of what we both know about our problem, and we both know this. I think that you will think I will choose the less pleasant place, because you know I would rather meet you there than not meet you at all, and you know I think there is a good chance I will miss you if I go to one of the places that is not distinctive, and the only other alternative is the less pleasant place. Since I think you think this, I think you will go to the less pleasant place. Therefore I will go to the less pleasant place." And we will both go to the less pleasant place, thus achieving co-ordination.

Certain things should be noticed about solutions to co-ordination problems achieved through salience. There is a good deal more "problem solving" involved than there is in cases of co-ordination achieved by agreement. In cases of co-ordination via salience, hypotheses about the state of another person's mind such as what he is likely to notice about a situation, how intelligent he may be, and what he is likely to believe in these regards about the other persons involved, serve as a basis for conclusions about the expectations he might have about what persons in the situation might do, and thereby help us decide what he will do. This is a property that is shared by our deerhunt example above. This property is not shared by co-ordination situations in which there is strong coincidence of interest, when they are solved by agreement. When an agreement exists, the fact of strong co-incidence of interest makes it

unnecessary for the persons involved to infer what the other persons involved will do; the agreement and the nature of the situation make it apparent what they will do. (Unless, of course, other aspects of the situation interfere. For example, if there is reason to expect that some person involved is mentally retarded or insane, and will not remember the agreement or thought they were talking in code when they made it, the other persons involved will again have to come to co-ordination by solving a problem. The point here is that under normal circumstances it makes perfect sense to say that an agreement does not so much enable a solution to be found in a co-ordination situation as it provides the solution to the problem.)

The last way that Lewis considers in which solutions can be found in co-ordination situations is via precedent. It may be that the persons involved in a co-ordination situation have been in a similar situation before and have solved it in a particular way. Or, they may have heard of a similar situation that has been solved in a particular manner in the past. This will give both a reason to recognize a possible solution, given their common background information. Thus, our telephone callers may have heard of people solving the problem they are in by the person who placed the original call calling back and the person who was originally called waiting until called. They may recognize the precedent as being the only alternative likely to be chosen by both, on the basis of its being the only aspect of

either's background relevant to their problem, and act similarly.

Because a precedent may have many possible interpretations, complex hypotheses about what the other persons involved will recognize as the salient aspect of a prior solution may be required for successful actopm on a co-ordination solution. For example, you and I may be in the position of our telephone callers, and be aware of a situation wherein the problem has been solved by the person who placed the original call calling back. We may also be aware that the person who called back was the shortest of the two, makes the most money, and had the birthday closest to the day of the call. A decision as to what aspect of the precedent is likely to be noticed and be considered the most salient and thus serve as a reference point enabling solution may have to be made by the persons involved.

When a co-ordination situation regularly occurs and is responded to in uniform fashion by the people among whom it occurs, such a regularity stands as a pervasive precedent for new occurrences of the problem. Though the people involved in an instance of such a situation may not have encountered the situation before with the persons presently involved, they may be sure that the others are familiar with the situation and its usual method of solution. When such a regularity exists, we have the basic structure required for there to be a convention operating among the persons regularly involved in the recurring co-ordination situation.

There are differences between successful co-ordination due to the existence of a regularity (i.e. pervasive precedents) and co-ordination due to salience (perhaps salience due to few precedents) which are similar to the differences pointed out above between co-ordination reached by agreement in co-ordination situations where there is strong co-incidence of interest and normal circumstances, and co-ordination reached through problem-solving via salience. When there are few precedents available for a co-ordination situation, it may be necessary for the persons involved to attempt to replicate the reasoning of the other persons involved through complex hypothesis about the minds of the other persons involved. ("How intelligent are they? What aspects of the precedents will be taken by them to be salient? How intelligent do they think I am? What aspects of the precedents will they expect me to take to be salient? Expect me to expect them to be salient?" And so on, perhaps.) As the number of precedents increases, the number of possible aspects of the precedents that can be interpreted as operative in the solution decreases (C, p.40), and the likelihood that others involved are unfamiliar with the situation and its normal mode of solution drops to nil. As such, the need for inference as what the others involved will do decreases with the ubiquity of the precedent. Unless there is some reason to expect that the instance of the co-ordination situation is not normal, and the others may not act as they can be expected to, we

need not take a problem-solving attitude to instances of the situation. That is to say, a regularity in handling a co-ordination situation does not enable us to solve a problem that confronts us, so much as it is, in fact, the solution.

Lewis's discussion of our methods of reaching co-ordination in situations in which there already are (existent) regularities of action does not mark the inadequacies of a problem-solving model. He treats precedents for co-ordination as properties of the situation that permit a particular co-ordination to be salient to the persons involved (C, p.36, p.38). He then goes on to treat regularities with respect to a recurring co-ordination situation as cases in which a large number of precedents permit the persons involved to co-ordinate their actions in essentially the same way that situations with a single precedent permit co-ordination. It is not that a regularity permits us to take for granted a method of coping with a co-ordination situation, but that the regularity permits us to generate the expectations about the state of other's minds characteristic of co-ordination problems solved by salience: "Since our present problem is suitably analogous to the precedents, we can reach a co-ordination equilibrium by conforming to this same regularity....If we do conform, the explanation has the familiar pattern: we tend to follow precedent, given no particular reason to do anything else, we expect that tendency in each other; we expect each other

to expect it; and so on. We have our concordant first- and higher- order expectations, and they enable us to reach a co-ordination equilibrium" (emphasis mine, C, p.39). That is, the conclusions about other's minds characteristic of co-ordination situations in which we must take a problem-solving stance in order to achieve co-ordination (the expectations about other's expectations and so on) are seen by Lewis to be essential to our reaching co-ordination in situations wherein there are regularities of action available. We do not, according to Lewis, develop concordant expectations about what will be done and act upon them. We develop concordant expectations about what is characteristically done, come to conclusions about the state of mind of other's involved, and then act in accordance with the regularity. We do not merely "tend to follow precedent, given no particular reason to do anything else", according to Lewis. I think he is wrong about this. His being wrong about this leads to problems in the analysis of convention that he develops.

The distinction I want to insist on can perhaps be focused by considering the book that inspired Lewis's application of co-ordination problems to convention, Thomas Schelling's The Strategy of Conflict (C, p.3).³ Schelling considers co-ordination problems in order to throw light on problems involved in the tactical use of nuclear weapons.

Roughly, the situation is that the only salient

³Thomas C. Schelling, The Strategy of Conflict (Cambridge, M.A.: Harvard University Press, 1960).

co-ordination is for neither side to use them at all. It is difficult for either adversary to distinguish other possible co-ordinations (e.g. use them only against military targets involved in the other side's ability to conduct an otherwise conventional war) from a decision to go for the big win (i.e. to blow up only part of the world). Thus Schelling's emphasis is on a case wherein reaching co-ordination is possible (a) in general, only via salience and (b) under the spectre of unilateral deviation. Such a case is significantly unlike a large number of the cases in which we solve co-ordination problems. In this context, Schelling reports research that demonstrates that persons involved in novel co-ordination situations are able to achieve a "meeting of minds" (Schelling's happy choice of phrase) in the absence of communication. It is clear that the persons involved in such a problem situation would have to depend on inferences and conjectures about the others involved that would not be essential to persons in co-ordination situations wherein a "meeting of minds" had already been effected and was not in question, that is, where there had been an agreement or there was a convention as to what was to be done.

In this section I have considered co-ordination situations and their methods of solution. I have suggested that a problem-solving model sometimes gives an inappropriate description of our ways of coming to co-ordination. When there is an agreement, or when there is

a regularity with respect to a co-ordination situation in the population in which it occurs, the need for inference to conclusions as to what the others in the situation will do is obviated, unless the case is unusual in some way. Two ways in which the co-ordination situations can be unusual in this sense have been pointed out. If there is only limited co-incidence of interest in the case, questions may be raised as to the possibility of unilateral deviation from an agreement or regularity. In such a case, co-ordination may depend on such things as beliefs about the state of mind of the other persons involved, beliefs about their beliefs about one's own state of mind, and so on. Also, if there is reason to believe that the others are not likely to act in accordance with an agreement or are not apprised of the nature of a regularity in coping with the co-ordination situation, similar considerations come into play in reaching co-ordination.

II. Lewis's Problem-Solving Model of Conventional Activity

David Lewis's most recent published version of his analysis of convention occurs in 'Languages and Language' (LL). Because the aspect of his analysis that I wish to attack is most easily presented with respect to this version of his analysis, I will treat of it first. After describing Lewis's analysis and presenting a counterexample derived from the concerns of the previous section, I will attempt a diagnosis and repair of Lewis's analysis.

Lewis's definition of convention, as given in 'Languages and Language' is as follows:

A regularity R, in action, or in action and belief is a convention if and only if, within P the following six conditions hold. (Or at least they almost hold. A few exceptions to the "everyone"s can be tolerated.)

- (1) Everyone conforms to R
- (2) Everyone believes that the others conform to R
- (3) This belief that the others conform to R gives everyone a good and decisive reason to conform to R himself.
- (4) There is a general preference for general conformity to R rather than slightly less than general conformity--in particular, rather than conformity by all but any one.
- (5) R is not the only possible regularity meeting the last two conditions.
- (6) Finally, the various facts listed in conditions (1) to (5) are matters of common (or mutual) knowledge: they are known to everyone, it is known to everyone that they are known to everyone, and so on. (LL, pp.5,6)

'David K. Lewis, 'Languages and Language' in Minnesota Studies in the Philosophy of Science, Vol. VII, ed. by Keith Gunderson (Minneapolis: University of Minnesota Press, 1975).

In order to see how this definition works let us consider a convention that developed in Oberlin Ohio, where it was once the case that all telephone calls were cut off after three minutes. It came to be that when this happened, residents would call back if they had been the original caller, and wait if not. Everyone conformed to this regularity and everyone believed everyone else did. The belief that everyone else acted in accordance with the regularity gave everyone good and decisive reason to do the same, and they all would prefer that no small number did otherwise. (If nearly all of them did otherwise that might be fine, if all did otherwise in a co-ordinated manner. Only a few of them doing otherwise would create problems, and they would prefer to avoid this.) And they could have, after all, had the person receiving the original call call back, and had expectations and preferences consistent with such a regularity.² Finally, the existence of the regularity is a pervasive precedent permitting them to know that all know the above, know that all knew it, and so on.

Consider now the following case, a population wherein everyone drives on the right. Everyone in the population has only ever seen the others drive on the right. Everyone has the strongest preference to drive on the right if the others do, and each hopes that no one gets it in their head to

²This alternative would not work as well: The person called would not always know the number he should call. It would, however, work better than no convention at all. Perhaps it would come to be the polite thing to do to state the number one was calling from at the beginning of a conversation.

drive on the left, unless, of course everyone gets it in their head to drive on the left. Everyone knows these things, and everyone knows everyone knows these things, and so on. So far, so good; by Lewis's analysis these people are parties to a convention of driving on the right. They could be ourselves. These people are, however, different from us in an important respect. Their driving on the right depends on their recognition of the other's intention to drive on the right. While driving, they hug the outer right line of the road, in order to ensure that their intention to drive on the right is recognized; if anyone strays from this position, those who see him do so immediately pull over and park, watching carefully to ascertain whether it is a lapse of attention or an expression of an intention to drive on the left. While on the highway, it is only rarely that anyone attempts to pass a slower vehicle, for they think that if oncoming traffic were to see them on the left, chaos would reign. The more intelligent among them refuse to drive at all, for they realize that not only must intentions to drive on the right be mutually recognized, but that intentions to have their intentions to drive on the right recognized must also be mutually recognized, and they cannot see how this might be managed. Stress-related illness and suicide are common among those who must drive for a living.

I think it is clear that these people do not have a convention going for them. The reason why they don't is that they are treating what is for us conventional as a problem

in need of solution. It is part of what a convention is, that it permits us to automate aspects of our living together and thereby get on with the business of life. Lewis's analysis does not require this, and it should.

Before suggesting ways in which Lewis's analysis might be repaired in light of the counterexample above, I'd like to consider the ways in which Lewis's problem-solving model of co-ordination insinuates itself in his analysis of convention. The aspect of his analysis which stems from his recognition that some co-ordination problems are solved in virtue of persons coming to concordant first- and higher-order expectations via hypotheses about other's minds is the requirement of common knowledge. Lewis does not intend that this requirement imply that persons involved in a convention have reasoned to that knowledge, i.e. have treated the convention as if it were a problem. As he says (LL, p.6) "The knowledge mentioned here may be merely potential: knowledge that would be available if one bothered to think hard enough." This suggests that the common knowledge requirement is supposed to mark reasoning that could be carried out, but which in general is not. It seems to be intended as an articulation of the justifications we might give if we were to consider in a calm moment our conformity to our conventions, though not necessarily the process by which we come to conform in a particular case. This interpretation is supported by Lewis's remarks (LL, p.25f) when he addresses the question of how an action can

be habitual yet "rational": "A habit may be under the agent's rational control in this sense: if that habit ever ceased to serve the agent's desires according to his beliefs, it would at once be overridden and corrected by conscious reasoning." In a similar context (C, p.141), Lewis says "[We] do not have to consider the agents' actual reasoning. We have to consider only the rational justifications of their choices by practical reasoning they could go through, given their beliefs and desires. Yet this is not to renounce an interest in explaining their choices. Justifications do explain choices, whether or not the agent actually goes through a process of reasoning following the justification. For it is a fact of human nature that we tend to act in ways justified by our beliefs and desires, even when we do not think through the justification. I may put it negatively: whatever may be the habitual processes that actually do control our choices, if they started tending to go against our beliefs and desires they soon would be overridden, corrected and retrained by explicit practical reasoning" (C, p.141). This 'agent's justification though not necessarily their reasoning' interpretation of common knowledge demonstrates Lewis's reliance on a problem-solving model in ways other than its acceptance of 'possible their reasoning'.

In saying that "justifications do explain choices" in this context, Lewis commits himself to the view that the relevant justification, if it were to be given, would

include conclusions of the form described in the definition of common knowledge (e.g. I know that he knows I will conform to R). Lewis makes this explicit (LL, p.6) where he claims that "This condition [common knowledge] ensures stability. If anyone tries to replicate another's reasoning, perhaps including the other's replication of his own reasoning, ..., the result will tend to reinforce rather than subvert his expectation of conformity to R." But consider: if the other's conformity is of the 'habitual yet rational' type Lewis describes, Lewis can't mean the above, as the reasoning that is supposedly being replicated doesn't occur. What Lewis should say is something like "If anyone tries to replicate the reasoning another might do if the convention were to stop serving their interests, this replication would reinforce rather than subvert his expectation of conformity." But this is very odd. If a regularity stops serving our interests, consideration of how others will reason to decide what to do will not likely reinforce our expectation conformity to that regularity, because by hypothesis the regularity is no longer furthering our interests. In such a situation, we might be entirely baffled, and not know what to expect of the other.

Justifications resulting in common knowledge are characteristic of situations that are specifically not conventions. How do we justify action in accordance with a convention? Consider driving on the right hand side of the road. In the population in which I live, such a practice is

conventional. When I drive on the right, do I believe that drivers in the oncoming traffic will expect me to expect him to stay on the right? I do, though I might have to think about it before I say so. (This is a potential belief.) Such a belief on my part does not, however, have very much to do with my reasons for driving on the right. I drive on the right because we always do, and I have no reason, in almost all instances, to expect anything else from anyone else. Conventions are, in general, cases wherein co-ordination is achieved, not through hypothesis and beliefs I might have about the beliefs and expectations of others in the situation in which we find ourselves, but through zeroth-order expectations that are justified directly by the conventionality of our co-ordination. (Lewis suggests (C, p.32), that ordinary situations can justify lower-order order expectations directly, without recourse to higher-order expectations. Conventions are precisely those co-ordination situations that are very ordinary.) In situations where, say, my belief that others recognize my intention to conform to a convention is part of my reason for conforming to the convention, it seems that there must be something abnormal about the cases. Higher-order expectations come into play only when there are characteristics of the situation that suggest the situation need be solved by such a mechanism. Perhaps I am driving in a crowded parking lot, where the rules of the road are not being strictly observed, and I must pay careful attention to

what the other drivers are doing, rather than expect them to keep right as a matter of course. One thing I am sure of: if attempting to justify my driving on the right on the highway, I certainly wouldn't be concerned to describe what I might think in order to decide what to do in the event our convention of driving on the right was deteriorating before my eyes. It would be irrelevant. Why would anyone justify what he does in one situation by giving a justification of a course of action in a completely different situation?

Lewis intends the inferences resulting in common knowledge (in the sense in which these inferences need not explicitly be made) to explain why conventions are stable (LL, p.6); the higher-order expectations that can be derived are "involved in sustaining conventions" (C, p.56). If such inferences and expectations do not describe actual reasoning, and are not, as a matter of fact, involved in the justification we might give of our conventional behavior, I do not see how these claims can be true. I have difficulty seeing them as other than unnecessary and misleading imports from his discussion of co-ordinations achieved via problem-solving.

I want to claim the following: if an action is conventional it cannot be a product of the kind of practical reasoning implied by Lewis's definition of common knowledge, and that our conventional activity is justified and sustained by our ability to recognize cases of the co-ordination situation they solve and act in accordance

with our convention, thus furthering our interests. That is, our conventions are sustained merely by their not breaking down, or giving evidence of their possibly breaking down.

Lewis has given two other versions of what common knowledge, both of which come nearer the mark than does the 'Languages and Language' version. In Convention, he gives a definition that is stated entirely in terms of the evidence that is available to a population in which a convention exists (C, pp.60-68). This evidence is essentially the evidence that is required to generate the beliefs implied by the 'Languages and Language' definition. In its essentials, it requires that there be no reason to believe the convention might not be observed. This, in itself, sounds much like my claim that the situation must be able to be perceived by those involved as a normal instance of the kind of situation the convention solves. This is not, however, the use Lewis intends, as demonstrated by his emphasis on the nature of the beliefs generated (C, pp.60-68), the way in which they are generated (C, pp.53-58), and his claim that these beliefs are "involved in sustaining conventions" (C, p.56). In these regards, this definition suffers from difficulties identical to those discussed above.

In what is apparently an aside, Lewis puts forth a third possible definition of common knowledge, one which seems to me to demonstrate that Lewis's own intuitions are not entirely inimicable to my own. He says (LL, p.6) "Perhaps a negative version of (6) would do the job: no one

disbelieves (1) to (5) hold, no one believes that others believe this, and so on." This is significant in that it cannot be interpreted as a justification of action in accordance with convention except of the "everything is apparently going normally"-type I favour. A justification of an action can be given by describing the beliefs or inferences of the person doing the acting (including here, his beliefs or inferences he might have should justification be required). This definition of common knowledge does not do that. It states certain beliefs are not had by the person involved. The only way this kind of requirement will justify action in accordance with a convention, or serve to explicate the stability of a convention, is if action in accordance with the convention is *prima facie* justified, possibly overridden by the kinds of considerations this definition rules out. Why do we act in accordance with convention? Because we always do, and negative common knowledge seems to hold. Why does the convention continue? Because negative common knowledge generally holds; our action in accordance with the convention does not tend to get brought into question by the kinds of concern negative common knowledge rules out. While it would be wrong to put too much weight on Lewis's suggestion of this possible interpretation of common knowledge, I think his acceptance of it as a candidate suggests that the concerns I've raised to this point apply to his concept of convention as well as my own, and an attempt should be made to remedy apparent

flaws in his analysis.

The alteration to Lewis's analysis I suggest is relatively simple. Lewis's condition (3), the condition that requires that everyone in a population is given good and decisive reason to conform by their belief that everyone else conforms to R, can be altered by adding the following proviso: If any other belief about the others involved is essential to anyone's conforming to R, then the regularity does not operate as a convention in that instance. This is not to say the regularity cannot be a convention; a convention can exist, and one can act in accordance with it without availing oneself of its conventionality. One can, in certain situations, have come to a decision that what is conventionally done is the best thing to do at the time.

This alteration to Lewis's analysis blocks the counterexample with which I began this section (i.e. the people who drive on the right in virtue of their recognition of others intention to drive on the right). If an action is taken on the basis of the kind of practical reasoning characteristic of co-ordination achieved via problem-solving, beliefs other than the belief that all others involved conform to the convention are essential to conformity, and the action is not taken in virtue of its being a conventional solution to the co-ordination situation. If, on the other hand, persons conform because everyone else does, and there is no reason to think they will do otherwise in the particular case, the action is

action taken in virtue of its conventionality. In the counterexample to Lewis's analysis no one in the population coordinates his action with those of the others in virtue of a convention of driving on the right. (Can one have a convention when no one avails themselves of it? It seems one can. Consider our manner of dealing with whose turn to go it is at a four-way stop. The regularity of proceeding in the order in which the drivers reach the intersection is intuitively a convention, as it is by Lewis's analysis. Only rarely, however, does anyone merely assume the others will follow the convention as a matter of course. (e.g. people proceed slowly when initially entering the intersection in order that their intention to proceed may be recognized.) What can be said about this case? A number of factors, the negative consequences of driving into someone else and the high proportion of persons breaking the convention relative to these negative consequences makes it pay to not avail oneself of the convention. The nature of the co-ordination situation and its history seem to be similar to the deerhunt example described in Section One. However one decides to describe the case (e.g. there is a convention, but persons do not avail themselves of it, or only partially, or something), the differences between this case and other conventions are able to be marked by my suggested alteration to Lewis's analysis and cannot apparently be marked by Lewis's analysis. This seems to speak for my version. The case is not just a convention.)

My suggested alteration to Lewis's analysis has a second positive consequence. In Convention, one of the uses to which the definition given there of common knowledge is put is to rule out a certain class of cases from being called conventional. These cases are those in which persons in the population within which what would otherwise be a convention have certain false beliefs about their conformity. Perhaps, for example (C, p.59), '[everyone] holds this false belief (call it f): "Except for myself, everyone drives on the right by habit, for no reason, and would go on driving on the right no matter what he expected others to do."'

Two things must be noted about Lewis's attempt to rule out this kind of case. First, he presents it as if it is ruled out by the definition of common knowledge given in Convention. It is not, as this definition is given only in terms of the evidence had by the population. While this evidence may make this belief unlikely, it does not make it impossible; it certainly does not logically preclude it, which is what is essential if it is to rule the case out by definition. Second, Lewis's concern with this case further illustrates his commitment to the problem-solving model of convention. Why would someone be concerned to rule out the case described above? Only, it seems to me, if they believed that a person's asserted beliefs about a convention necessarily reflected his reasons for conforming, that is, if conventional activity were best understood as a

consequence of practical reasoning.

Lewis's rejecting regularities about which we have false beliefs as being conventional strikes me as heavy-handed when those beliefs are not part of our reasons for conforming. In the population which Lewis describes the false belief they hold does not constitute part of their reason for driving on the right--at least, this does not strike me as the most natural way to interpret the example--they drive on the right because they always drive on the right and there is no reason to believe anyone will stop now. That is, they drive on the right for the same reasons we do. If it strikes the reader as odd to say that our discovering something about why we act according to a certain regularity is not a matter of our discovering why something is a convention, but is rather a matter of something that was not a convention suddenly becoming one, he should agree with my concerns on this. Beliefs we have about our conventional activity need only be seen as affecting whether or not a regularity is a convention if those beliefs are our reasons for so acting. My position permits, and Lewis's apparently does not, our having innocuous false beliefs about our conventional activity. Application of my suggested change to Lewis's clause (3) gives a better resolution to Lewis's concerns about false beliefs that occur in a population. If those false beliefs are essential to their conformity, then we do not have a convention. If they are not, then we do. This line of

argument might seem to demand alteration to Lewis's condition (6) in his definition of convention, his requirement of common knowledge. It does not, if we construe "common knowledge" strictly according to Lewis's definition in Convention. That definition puts a condition on the evidence that must be available to a population for there to be a convention operating in it. (In this regard, the definition of common knowledge can be taken as an articulation of what must be the case for there to be appropriately said to be a regularity in a population.) We can preserve that definition, let those higher order expectations and beliefs as are generated in a community via mutual ascription of rationality be generated as they may, and permit my suggested proviso to condition (3) to eliminate those causes wherein strange beliefs function as part of person's reasons for conforming to what might otherwise be conventional from being properly called conventional. On this account, populations such as the one in Lewis's example would be very unlikely, although possible, conventions.

I take the results of this section to be as follows. It has been shown that our conformity to a convention cannot be action taken in virtue of that convention if it is accompanied by practical reasoning of the kind suggested by Lewis's analysis. Also, it has been argued that higher-order beliefs and expectations do not serve to justify our action in accordance with convention and thereby account for the

stability of our conventional solutions to co-ordination problems. To the contrary: conventions are stable because they further our interests and we have little reason to believe they will break down. That we are able to follow them unreflectively is both an essential ingredient of our concept of convention and a consequence of the reason why they are stable.

III. Behavioral Evidence for Higher-Order Reasons Underlying Conventions

David Lewis's account of our epistemic relation to our conventions relies on the ordinary senses of the words "intention" and "belief", and the ordinary criteria for their application. Jonathan Bennett, in his Linguistic Behavior (LB), has attempted to develop an account of intention and belief that relies solely on possible observations of the organisms to which we might wish to make these attributions.' Before considering the conventionality of language and conventional meaning, the topics that will occupy my attention for the rest of this thesis, I wish to consider and reject the possibility that behaviorally-ascribed higher-order intentions figure in our conventional activity in any non-trivial manner.

In the present context, I take a particular attribution's purported role in governing our to be trivial if it can be replaced without loss with a negative formulation. That is, in order for a belief that \emptyset to non-trivially govern our behavior, there must be some reason to choose to attribute it over the absence of a belief than not- \emptyset . The view of convention-governed behavior I advocate, claiming that we are apprised of how to act in virtue of our ability to recognize instances of a convention-governed situation when there is an absence of reason to believe the convention is not operating normally, would be contradicted

'Jonathan Bennett, Linguistic Behavior (Cambridge: Cambridge University Press, 1976).

if non-trivial attributions of higher-order propositional attitudes could be defended.

Bennett considers direct and indirect evidence that can be obtained to show that a regularity in a population is followed as a convention. As he accepts Lewis's analysis of convention, with the 'Languages and Language' version of common knowledge, this involves showing that higher-order intentions figure in our conventional action. (I will call a regularity with the properties Lewis's analysis suggests, with the 'Languages and Language' version of common knowledge, a Lewis-convention.)

Bennett's "direct evidence for conventionality" is not to the point (LB, pp.181ff). He proposes that in order to ascertain the reasons for which persons act in accordance with conventions, say, by doing Ψ in situation \emptyset , we examine "cases where we can learn that someone's believing Q is a reason for his doing Ψ because we find he doesn't do Ψ when he doesn't believe Q " (LB, p.183). Of course, this is not quite what he means. Rather, he means we should consider cases where we can learn that someone's believing Q is a reason for his doing Ψ because we find he doesn't do Ψ when he believes the denial of Q . His position is better expressed in "[any] looped belief [i.e. belief about another's belief about his own beliefs] of x 's which is evidence that for him the regularity is a convention must be a belief that credits someone else with what x believes to be a mistake" (LB, p.184). In order to find out the beliefs

in virtue of which a person acts in accordance with a convention in a \emptyset situation where he has no reason to think that, for example, the others involved think he intends to do something other than Ψ , we are to examine what he does when he believes that others believe he will not do Ψ .

I find this odd. In a situation where someone believes something that suggests a convention will not operate properly, he will evince behavior that speaks to his ability to recognize relationships between the beliefs he attributes to others and what actions they will likely take, and the relationships between certain actions taken and success in the co-ordination situation at hand. We can find out if he can solve a co-ordination problem for which there is conflicting evidence for how people will act. In order to claim, as Bennett wants to, that the beliefs and inferences that can be attributed to a person in such a situation are those in virtue of which he conforms to convention under normal circumstances, more is needed. If someone doesn't do Ψ when he believes not- Q , all we may infer is that he doesn't believe not- Q when he does do Ψ , and that this is essential to his Ψ -ing. To claim that he believes that Q is also fair enough, but the situations Bennett describes do not support a claim that he Ψ s in virtue of a belief that Q in any way that distinguishes it from the mere absence of a belief that not- Q . Bennett's direct evidence supports only a trivial attribution of Lewis-conventionality. Without the intentional camouflage, what can be supported by Bennett's

direct evidence is a version of negative common knowledge. And this, we have seen, shows marked affinities to the "no essential, higher-order propositional attitudes" point of view I espouse.

Bennett's direct evidence for claiming someone acts in accordance with a convention virtue of higher-order propositional attitudes is open to another difficulty. Should anyone ever carry out his program of observation and research, when they come to search for evidence for higher-order propositional attitudes underlying conventional behavior they could very well find a person's behavior to give only ambiguous evidence for the "looped thoughts" Bennett seeks. If, faced with a situation in which he believes the others believe he will not conform to a convention for whatever reason, a person were to "panic, or be overcome by bewilderment at the ability of the situation, or let the moment for doing Ψ slip by while he mused about what to do (LB, p.191), the direct evidence the researcher seeks higher-order intentions might not be forthcoming. Bennett's position seems to be: if we ever evince higher-order propositional attitudes in solving abnormal cases of conventionally solved co-ordination problems, then this licenses our claiming all conventional behavior operates in virtue of such things. Instances of similar situations where we fail to find justification for such attributions, on the other hand, do not correspondingly give reason to claim that such inferences, etc. do not occur in

our conventional behavior, or only sometimes occur, or are not essential to it. Looped intentions, are on this account truly strange and magical beasts, not obviously related to person's behavior, contrary to Bennett's stated objectives.

Bennett's indirect evidence for Lewis-conventionality does not fare any better. Roughly, indirect evidence for higher-order beliefs is justified via an attribution of general problem-solving abilities to the persons involved in (what is supposed to be) a Lewis-convention and attributing to them ever more complicated beliefs on the basis of the evidence we know they have. This amounts to declaring that a person's practical reasoning includes as premises nearly anything which would justify their action in accordance with a convention, thus licensing such conclusions as x does Ψ because he thinks that y thinks that x expects y to do Ψ .²

²Bennett's indirect evidence is essentially the third person use of the attributions of rationality described by Lewis in order to justify higher-order attributions within the population. If we have evidence that a person's reasons are transitive, that is, if we have evidence that the persons we are considering generally "realize that if P 's truth is a reason for doing X , and that Q 's truth is a reason for P 's truth, then Q 's truth is a (deeper) reason for doing X " (LB, p.188), we have a reason to say that those persons have Q as a reason for doing X when we have reason to believe they believe that Q . Because we can have evidence that, in a \emptyset -situation, x does Ψ because he expects y to do Ψ and y does Ψ because he expects x to do Ψ , we know that (assuming an ability on x 's part to "handle evidence properly" we have "general and admittedly inconclusive grounds" (LB, p.189) for believing that x thinks that y will do Ψ because he thinks x will do Ψ . Thus, given a general premise about the transitivity of x 's reasons we have evidence that x does Ψ because he thinks that y expects him to do Ψ . Via a similar inference--based on our last inference and its "general and admittedly inconclusive grounds", a premise about the general ability of x to handle

Two points need be made. First, that someone realizes in some sense that a certain fact is a reason for doing something does not entail that that fact is one of his reasons for doing that thing. A premise about a person's problem-solving abilities supports conclusions about his ability to realize that certain evidential relations hold. More is required to show that it is in virtue of the attributed problem-solving ability, the evidence he has, and the conclusions that are possible from this evidence, that someone acts according to a convention: it must be shown that the problem-solving ability attributed is in fact being used. Indirect evidence of the kind Bennett suggests, then, will not distinguish between the hypotheses about the roots of conventional behavior that Bennett wishes it to, that is, between "convention" and "unthinking habit". (LB, p.182). His inability to do so supports my claim that the dichotomy is a false one.

The problems for Bennett's indirect evidence do not stop there. By his own admission, it amounts to pretty weak grounds for the attributions he wants to make on the basis of it. It is not even strong enough to imply that the direct evidence considered above would also be available:

Our basis for [higher-order attributions] is so indirect and thin that it would be risky for us to conclude that if x were in a \emptyset situation where he did not think that y thought that x expected y to do Ψ , he wouldn't do Ψ in that situation...for all we know, he would panic, or be overcome with

²(cont'd)evidence, and the transitivity of x's reasons--we may say that x does because he thinks that y thinks that x expects y to do .

bewilderment at the oddity of the situation, or let the moment for doing Ψ slip by while he mused what to do. (LB, p.191)

This leads Bennett directly into the problem discussed above: a negative result in an attempt to find direct evidence for essential higher-order propositional attitudes in conventional behavior should support the null hypothesis (that is, that they are not essential), if we are prepared to accept a positive result as supporting a claim that they are essential. Also, Bennett considers the possible objection that "'the thin construal of [higher-order reasons] you are suggesting is an abuse of the words'", and responds "that is not unreasonable" (LB, p.192). He goes on to defend such a construal on the grounds that it is necessary, on pain of finding that what we normally take to be conventions are not, as Lewis's analysis of the concept of convention is the only one we have: "The only alternative would be to drop Lewis's account of what a convention is; but what other account is there?" (LB, p.192). I have suggested that the aspect of Lewis's analysis which leads Bennett to concern himself with higher-order reasons is a product of inappropriate application of a problem-solving model of conventional behavior, that what we normally take to be conventions do not depend for their being so on such things. Thus, there is no theoretical or conceptual justification for Bennett to accept substandard behavioral evidence with respect to the role of higher-order propositional attitudes in conventional behavior.

Bennett's failure to describe how there could be behavioral evidence that there was a Lewis-conbention operating in a population has two important consequences. First, it supports my claim that we do not come to co-ordinate our actions in situations governed by conventions in virtue of our having concordant higher-order expectations about our behavior. Consider: evidence about our higher-order expectations that is unavailable to Bennett's observer is evidence that we ourselves lack. If concordant higher-order expectations are claimed to figure in our behavior in situations governed by conventions, it must also be admitted that these higher-order expectations are not generated by behavioral evidence that is available to us. Conventions must then be thought to depend on a set of baseless beliefs. I do not think anyone wishes to embrace this consequence. Further: what justification we have for attributions of higher-order propositional attitudes must flow from the fact that we know how we can be expected to act in a situation governed by convention, to what higher-order expectations we may have about our behavior, and not the other way around. Thus, our conventional behavior cannot be thought to be governed by our coming to those higher-order expectations.³

Second, the attempt to show a behavioral basis for a claim that higher-order propositional attitudes are essential to conventional behavior is a kind of archimedian

³This argument owes much to John Biro's "Intentionalism in the Theory of Meaning", The Monist, vol.62, number 2, 1979.

point for the project of Bennett's Linguistic Behavior. Bennett applies Lewis's analysis of convention to the phenomenon of conventional meaning, in order to support his claim that meaning can only be understood if it is considered in terms of the propositional attitudes of the speaker, and not, except in an eliminable manner, in terms of the normal, accepted, or conventional meaning of an utterance. Lewis's "reduction" of the operation of conventions to concordant expectations of action on the part of persons involved in a co-ordination situation not essentially different in kind from non-conventional co-ordinations, aided Bennett in his project. I take myself to have here given some justification for a more general thesis which is contrary to Bennett's intended conclusions: convention-nominalism is false. Any conventional activity is, as a matter of conceptual truth and for its continued operation, dependent upon the fact of there being normal cases, followed in virtue of their being normal cases. The cluster of problems relating to the conventionality of language and the nature of conventional meaning are the topics that will concern me for the remainder of this thesis.

IV. The Conventionality of Language

There is a significant distinction within the class of conventions that is obscured by Lewis's problem-solving model of conventional behavior that needs be marked. Some conventions, like driving on the right, or the original caller in Oberlin, Ohio calling back when the three minutes maximum was in effect, were once problems that had to be solved by people. We can take the same position with respect to the problems these practices solve as did the first persons that had to co-ordinate their actions. If one of these conventions did not exist, and we were confronted with the co-ordination situation that it solves we would have the capacities to solve it of our own efforts, either in the particular case or by the institution of a new convention.

Other conventions are not of this kind. Language did not emerge because people pretty much like us were confronted by a co-ordination problem and thought up a way to solve it. (Anything without a language is something that isn't much like us.) Since the time when whatever primordial communicative regularity our languages developed from emerged, it has gone through a process of development that leaves us unable to state clearly what co-ordination it achieves (What problem does it solve? Be specific.). Moreover, it leaves us largely unable to state the mechanisms by which it achieves it. In addition, the process of development by which our languages come to be as they are is intertwined with the social and genetic development of

the species in such a way that our availing ourselves of our linguistic conventions is not an intellectual accomplishment but a cognitive activity, dependent upon our having fairly specific genetic heritage and having grown up in a human society. That is, we could not, of our own efforts, create a solution to the problem of communication if we lacked it. Likewise, language is not the kind of thing that we learn because it is explained to us. Driving on the right is.

All conventions are not created equal. In the last section I articulated one way in which Lewis's problem-solving model of convention leads him to take up concerns that are misleading with respect to the reasons why we conform to conventional solutions. Some conventions, if they can be said to have emerged as solutions to problems at all, did so more in the sense that it true to say the feathers of birds emerged as a solution to a problem than in the sense that driving on the right emerged as a solution to a problem. In this section I will consider a way in which Lewis's discussion of convention is untrue to conventions of this type. (I refer to "conventions of this type" even though I cannot readily think of clear examples other than language. To the extent that we can be said to have ethical conventions, they would seem to be candidates. Legal traditions would seem to be similar in some respects to the linguistic case, as although they develop as solutions to particular problems, the net co-ordination that results need not be an intended result. I hazard a guess that many of the

complex social facts that cannot be adequately explained as the products of human design are roughly of this type.)

Before considering language's status as a convention that we could not introduce if we were to have to address it as a novel co-ordination situation, and its implications for an analysis of convention, a summary of Lewis's account of how language is conventional is appropriate (C, pp.177-95; LL, pp.7-12).

Not all of what can correctly be said to be "using English" can properly be said to co-ordinate speakers and hearers, that is to permit communication. The clearest example of such a thing is lying. Lewis's way of articulating precisely what the regularity is in a population that speaks a language reflects this. If a language *E* is used in a population, there exists a convention of truthfulness and trust in E. Speakers of *E* try not to utter false sentences; understanders of *E* impute trustworthiness to others and believe that sentences uttered in *E* are true, and hence come to believe what is said. (Lewis extends the usual definition of truth in order to take into account sentences that are not declarative: interrogatives, imperatives, commissives, etc. (C, pp.184-92; LL, pp. 13-15).) The regularity of truth and trust in a population seems to satisfy the definition of convention. There is a regularity of truth and trust in English in the population in which I live, and the people I associate with believe such a regularity exists. Not that

they have ever said so, or ever will; in any instance of communication they expect the speakers to be truthful and the hearers to be trusting and can tell whether, for a given utterance or belief on the basis of an utterance, whether or not it would be consistent with these expectations. That others act in accordance with truth and trust in English gives each of us good and decisive reason to do so too; if one of us is a speaker, we expect a hearer to be trusting, so we may as well tell the truth if we are to communicate (unless we are lying, but then we aren't acting in accordance with the convention. Touche.), if one of us is a hearer, and we expect truthfulness from a speaker we infer that the sentences of English he utters are true according to his beliefs, this by and large gives us reason to believe what he says is true. There is also a general preference for truthfulness and trust in English in the population. Given that most conform, everyone would rather all conformed: the odd untruthful speaker or untrusting believer, or truthful and trusting speaker of another language will disrupt communication. And an alternative to truthfulness and trust in English is possible and could be a convention were it a regularity, namely truthfulness and trust in some other language. And finally, ubiquitous past experience of truthfulness and trust serves as a basis for common knowledge of the preceding. (Consideration of the deerhunt example in section one above makes possible an articulation of what occurs when someone lies, and how it is this relates

to the existence of a convention of truth and trust in a language. Language serves, in general, cases of near-perfect co-incidence of interest. Sometimes, however, this is not the case. In certain cases, it serves the interests of one or more of the parties involved to defect from the prevailing co-ordination, that of telling the truth, by telling a lie. Likewise, a hearer can defect by not believing what he is told. It can be noted that the concerns raised in previous sections imply that if a hearer, before believing what is told him, considers and rejects the possibility that he is being lied to, he is not availing himself of the convention in that case. A higher order hypothesis is essential to his coming to the appropriate belief. (Cases wherein linguistic conventions serve as a background for novel communicative co-ordination problems will be discussed below.))

We have seen that it is necessary to eliminate as candidates for conventions those regularities that are followed for reasons besides that given in the third condition of Lewis's analysis. With respect to language, given the inaccessibility of adequate descriptions of the co-ordinations it permits us to achieve and how it achieves them, it is more than likely that a large number of our beliefs, even our considered beliefs, about our linguistic conventions and why it is we follow them are false. None of these beliefs, however, can plausibly be said to affect, in general, why we speak and believe what we do. Language

functions quite nicely in the face of a large degree of ignorance as to its nature. Parallel to the concluding arguments of the next-to-last section, we may say that what false beliefs we may have about language do not preclude our saying language is conventional, because those false beliefs do not constitute erroneous reasons for conforming to our convention. In certain kinds of cases, we can just be wrong about why it is we do a certain thing.

Because language is the kind of convention that was never approached as a problem, there are certain kinds of beliefs that seem possible in linguistic communities that Lewis's analysis rules out. One of the things required by Lewis's analysis is the common knowledge of the possibility of alternatives. Even on the strict interpretation of "common knowledge", which requires that there be an evidential basis for common knowledge of alternatives, although does not require any beliefs be had on that basis, it seems to me that this is too strong a demand.

Consider an isolated tribe in which neither in memory, legend or myth is there any acquaintance with other languages. In such a tribe, there is no basis for common knowledge that another language is possible. On Lewis's Convention version of "common knowledge", it cannot be said that their language use is conventional as there is no evidential basis for such knowledge. If people who have no evidence for the existence of alternative linguistic regularities have no beliefs that there are such things, as

is only likely, Lewis's 'Languages and Language' version of 'common knowledge' likewise implies their language use is not conventional. With minor adjustments to the example, even the weaker, negative, version of common knowledge implies these people's language use is not conventional; a belief which is clearly possible that is inconsistent with there being common knowledge of alternatives can be built into the example. Perhaps the tribe believe, as I am told primitives have believed, that there is a magical relation between an object and its name. Let us also say they do believe that this relation is not alterable. Thus they believe something inconsistent with, for example, the idea that objects could have different names than they do. They believe it is false that they could talk differently in this way, that people who tried to would be making errors.

It might be felt that, because of this tribe's isolation and peculiar beliefs, their language is not conventional. In this regard, I think Lewis's comment that it is "something only a philosopher would think of denying--that there are conventions of language" (LL, p.7) is appropriate. Taking it as axiomatic that language is conventional, the requirement that there be evidence for alternative possible regularities, or evidence and knowledge or alternative possible regularities, must be rejected. How exactly to go about rejecting it raises considerable difficulty.

Some conventions solve problems that were problems for the people involved because they were able to see alternatives in the courses of action that were open to them and able to see that their interests were only served if they could co-ordinate their actions with those of others involved. It is of the nature of conventions that solve problems like these, call them accessible problems, that the persons involved in them are aware of the problems they solve, and what alternatives there may be. Most everyone who has learned to drive on the right has had why he should drive so explained to him. For anyone else, the problem situation and its possible modes of solution are as obvious as they can conceivably be. In situations of this kind, I think we would be loathe to say there was a convention in a population wherein there was no awareness of alternatives, as this is essentially a lack of knowledge that there was a problem that had been solved. With respect to accessible problems, this amounts to a near incomprehensible blindness of the part of the population. Driving on the right just is not driving on the left, calling back just is not waiting. It is not clear how a problem could be recognized and solved without such awareness. For conventions which deal with accessible problems, Lewis's requirement of common knowledge seems appropriate.

Thus, there seems to be a distinction between co-ordination problems for which common knowledge of alternatives is an essential ingredient to there being

conventional solutions to them, and those for which this is not true. What difference does this distinction mark? Our description of an alternative to speaking English ("Speak Arabic!") does not enable us to begin observing the alternative described; we need access to someone who does speak Arabic, and time, and we need to be young enough that it has not become too difficult. Also, there is no relatively small set of alternatives that are defined by our being able to distinguish them, as is true with drive on the right/don't drive on the right. In the absence of English, a unilingual English speaker would be completely unable to achieve communicative co-ordination. (This deletes the possible effects of smiles, pointing and other more or less universal signals.) The distinction seems to be between those conventions that I've called accessible and those that are not, that is, those that are not solutions to problems, in the relevant sense of the word.

From the perspective of attempting to analyze "convention", this leads us into a problem. Common knowledge of alternatives seems essential to the conventionality of one class of conventions, and misleading with respect to the nature of certain others. One way of dealing with the problem would be to declare that language (and other conventions that are similar in this respect) are not conventions at all. I think that the cost of this would outweigh the gain. Language, in a significant number of ways is like driving on the right. Although in other ways it is

importantly dissimilar, to be unable to mark the similarity by a term that strongly suggests itself to us an appropriate to both (It is "a platitude that language is ruled by convention" (C, p.7)) would seem to me to ask something of our language that it will not give us. (A different name for every different thing, even when there is reason to consider them together.)

How then to deal with Lewis's requirement that there be common knowledge of alternatives? It seems to me best to leave it as is. We are aware of a way in which it misleads with respect to our relationship to our language, and this will serve to mark the sense of metaphor which is appropriate in calling language conventional.

One of Lewis's motivations in Convention was to articulate a concept of convention that makes it possible to say why language was conventional without resorting to metaphors. Prior to Lewis's work, the best approximation limited conventions to regularities that had been established by explicit agreement. This left anyone willing to call language conventional facing the obvious regress argument: language could not have originated by agreement, because agreement requires a previous language. Ergo, language is not conventional.' Lewis, with Schelling's demonstrations that persons can reach co-ordination in the absence of communication, was able to ward off the regress.

'This position is taken by W.V. Quine in "Truth by Convention" in his The Ways of Paradox and Other Essays (Cambridge, M.A.: Harvard University Press, 1976).

But what is wrong with the idea language originated in agreement is not that one can derive a vicious regress. The regress, once derived, shows something is wrong. The same thing can be wrong in an account of the conventionality of language that does not lead to such regress.

It seems to me that David Lewis's account of convention makes the same error as creates the regress in the agreement model of convention. The error is attempting to explain the existence and operation of conventions in terms appropriate to the solutions to problems reached by adult, linguistically capable human beings. We did not come to have a language, nor are we able to speak a language, in virtue of our being persons and having an understanding of the ways in which language furthers our interests.

The primordial way of explaining anything involves making reference to an object's intentions, beliefs, and desires, and telling a little story about it. With respect to complex things, this is sometimes the only way of explaining them we have. In Lewis's little story about our acting in accordance with conventions we are told we do so in virtue of standing in a relation to little problems. We know that this isn't the way it is with language. Hence we retain the sense of metaphor. (Not quite a metaphor: metaphors are false and obviously so. To call a language a convention is too close to the truth to be truly a metaphor. But: it is not just the truth.)

V. Conventional Meaning and Speaker's Intentions

David Lewis has proven (C pp. 152-59; LL, p.4) that his analysis of convention implies that conventional language use has the properties that H.P. Grice has claimed are essential to cases where someone can be said to mean something by what he says or does.¹ The emendment I have suggested to clause (3) of Lewis' analysis blocks Lewis' proof, and demonstrates an important difference between conventional meaning and cases of meaning that are adequately described by Grice's analysis. The purpose of this section will be to describe Grice's view of meaning and consider how it misrepresents conventional language use.² (In what follows I refer to utterances by which people mean something in virtue of there being communicative conventions as instances of "conventional meaning".)

Grice distinguishes between what he calls natural and non-natural senses of the word "meaning". Natural senses of meaning are those in which something is a natural sign of or evidence for something else, as in "Those spots mean measles". Non-natural senses are those in which it can be said that someone meant something by a certain act or utterance, or where it can be said that the act or utterance meant that p for some sentence, or where it can be inferred from a statement about what an act or utterance means a

¹H.P. Grice, "Meaning", The Philosophical Review, vol.66, No.3, 1957.

²There have been elaborations and alterations to Grice's theory by Grice and others. What is attacked in this section is common to all.

statement about what is meant by that act or utterance. Roughly non-natural senses of meaning are those which can be used to describe acts of communication.

Grice demarcates the difference between natural and non-natural meaning via consideration of cases similar to the following:

- (1) Herod leaves the head of St. John the Baptist where Salome will see it, supposing she will not realize that he has done so.
- (2) Herod presents Salome with the head of St. John the Baptist
- (3) Herod draws Salome a picture of St. John's severed head.

In all three cases Herod intends to get Salome to believe that St. John the Baptist has been executed. In the second and third cases, Herod intends Salome to recognize his intention that she should have this belief. In the third case, but not in the first or second, Herod intends that Salome's recognition of his intention that she should believe that St. John the Baptist had been executed to be part of her reason for so believing. A picture is not evidence, as a severed head is, that St. John the Baptist has been executed. It is only if Salome recognizes the intention with which the picture is drawn that she will come to believe that he has been executed.

Grice's intuition is that the third case above is a case of non-natural meaning and that the first two cases are

not. He takes what is essential for someone A to mean something by an utterance x to be a follows. A must intend to induce by x a belief in an audience, and he must intend his utterance to be recognized as so intended by his audience, and this recognition must be part of his audience's reason for coming to that belief.

Grice gives no argument to suggest that his analysis of non-natural meaning, based as it is on examples of non-natural meaning that are not conventional meaning, generalizes to cases of non-natural meaning that are also cases of conventional meaning. There is reason to question that it does. If, as in the example above, someone draws a picture in order to communicate something, his audience must rule out alternative possible interpretations of what he is attempting to do before they can come to the appropriate belief. They must not think he is doodling, or attempting to produce a work of art, or expressing a wish that the events depicted will come to pass, if they are to come to believe that the events have come to pass. In drawing a picture in order to communicate, it is essential, therefore, that the artist intend his audience to have their recognition of his intention be a significant reason of theirs for coming to believe what he intended. In a case such as

- (4) Herod says to Salome "John the Baptist has been executed".

however, alternative interpretations need not, in general, be ruled out. Why then, does Herod need to intend that

Salome have her recognition of his intention as part of her reason for believing St. John the Baptist has been executed? Why can't he just intend she believe what he tells her?

The argument in the literature which has provided justification for a Gricean diagnosis of conventional meaning is Lewis's. In Convention, a detailed argument is provided to show that in a conventional signalling system (e.g. one lantern means the British are coming by sea, two lanterns mean they are coming by land, etc.) the person giving the signals has Gricean intentions. There, he conjectures that any signal given in accordance with the general conventions of language will also demand a Gricean diagnosis (C, p.159). In 'Languages and Language' Lewis claims the Gricean mechanism operates in any convention, as each person's replication of the other person's reasoning in a co-ordination situation forms

part of his own reason for thinking and acting as he does.... X intends to bring about a response on the part of Y by getting Y to recognise that X intends to bring about that response; Y does recognise X's intention, and is thereby given some sort of reason to respond just as X intended him to. (LL, p.4)

I have argued above that, in normal conditions where there is a conventional solution to a co-ordination situation, higher-order attributions do not figure as part of the reasons persons involved in a convention have for their action in conformity to a convention. My emendment to Lewis's condition (3) specifically precludes this. That is, if the Gricean mechanism operates in a case of communication, then that communication cannot be said to be

a case of communication in virtue of a convention. Lewis's proof that conventional communication is Gricean depends on precisely that aspect of his analysis we have found to be untenable. Correction of that aspect of his analysis provides a demonstration of the denial of the statement that conventional communication is Gricean. If I intend your recognition of my intention to get you to believe that p to be part of your reason to believe p, then I am not availing myself of our convention of truth and trust in English. If I say something that means p in English, and you believe me, we conform to the convention, and I neither intend you to, nor do you, engage yourself in a problem of interpretation. We don't need to, for we have a convention that eliminates the problem.

I do not want to claim that, in the operation of a language, Gricean intentions never occur, but that those cases in which it occurs are those cases in which the conventions of the language are in question or are being manipulated in some way. One way in which this can occur has been suggested above. If I am talking to you, and there is good reason to believe I may be lying about what I say, (let us say I have much to gain by being untruthful), and I know we both know this, yet I wish to tell you the truth and to have you believe me, I may hope that you will recognize my intention to tell you the truth, and in virtue of your recognition of my intention, believe what I say. To make the case precisely Gricean, let us say that I tell you in the

course of our conversation that I am telling you the truth. I intend you to believe that I am telling you the truth in virtue of your recognition of my intention to tell you the truth. That is, I mean that I am telling you the truth without conventionally meaning that I am telling you the truth. (I can only conventionally mean that I am telling you the truth if I tell you this when I do not think you need to be told. If this seems to be a paradox, it is because one can act in accordance with a convention without availing oneself of its conventionality.)

I may be your friend and be engaged in a conversation with you about how such words as "friend" are misused by a large number of people, using them as they do to designate all kinds of relationships that are not friendships. I tell you that I am not your friend. Our mutual belief that the word has been misused to the point of making it useless for communicating what we think it should, might demand that my saying this to you, thereby intending to have you believe I am your friend, depends upon my intention that you recognize my intention to let you know that I am your friend. (Why else would I use a word that has been claimed has become meaningless? Perhaps to let you know that I am not really your friend. Perhaps I have put you in a position where you have to decide.) Again, I have meant that I am your friend without having availed myself of the conventional meaning of the words I use.

You may intend to let me know something in virtue of what Grice has called a conversational implicature.³ That is, you may fail to observe the limitations on what is appropriate in a conversation in such a way to let me know that you intended something quite different. Let us say (Grice's example) that I have asked you to write me a letter of recommendation for admission to a doctoral program. You write, and let me know that you have written, "Rod has good penmanship and always comes to class." Because what you have said is manifestly irrelevant to whether or not I am any good at philosophy, I realize that you intend the reader to believe that I am not any good at philosophy, and thereby come to believe that you do not think I am any good at philosophy. You have intended me to recognize your intention to let me know that you believe I am not good at philosophy, and thereby come to believe you believe this. Thus you have meant something by manipulating the normal expectations we have with respect to the conventional meaning of your utterance. (It is interesting to note that Grice demarcates conversational implicatures from what he calls conventional implicature by saying "The presence of a conversational implicature must be capable of being worked out; for even if it can in fact be intuitively grasped, unless the intuition is replaceable by an argument, the implicature (if present at all) will not count as a conversational implicature; it

³H.P. Grice, "Logic and Conversation" in The Logic of Grammar, ed. by Donald Davidson and Gilbert Harman (Encino, CA: Dickenson Publishing Company, Inc., 1975).

will be a conventional implicature." This difference, between communication that can be described as operating in virtue of inference from premises about the context (and what is done in that context) and conventional communication, is essentially that which is hidden by Grice's analysis of meaning.)

These examples of the Gricean mechanism operating are cases wherein a linguistic context has provided a communicative co-ordination situation, and co-ordination (the audience forming the intended belief) is achieved through the audience coming to a conclusion about the speaker's intention. In the next section I will consider certain phenomena that have been taken to demonstrate the operation of the Gricean mechanism within linguistic conventions. I will attempt to show that they do not.

VI. The Gricean Mechanism in Conventional Language Use

In the sixth chapter of Linguistic Behavior, Jonathan Bennett describes certain "familiar kinds of communicative transaction" which he claims "cannot be understood except in Gricean terms". The communicative transactions that he describes occur within the conventions of normal language use, and, if Bennett were correct about the necessity of a Gricean diagnosis of speaker's intentions with respect to these cases, a significant class of apparently conventional utterances would be counterexamples to a claim that language use is almost entirely non-Gricean. This, in itself, would not show my argument to this point to be invalid. It could be the case that in its normal functioning, language presented us with a class of cases that need to be treated as problems to be solved via lemmas about other person's intentions to have us recognize their intentions. If it is not the case that all or nearly all of language use functions in this way, such a situation can be assimilated to my argument to this point by saying that those aspects of language use depend on, but are not themselves instances of, conventions. But this is no matter, as Bennett is wrong in saying these aspects of language use are Gricean anyway.

The aspects of language use that Bennett thinks demand a Gricean diagnosis are of two kinds. The first are cases wherein there is an element of ambiguity in an utterance (LB, p.167ff). If a speaker utters something that has more than one possible meaning, or if his utterance can be

interpreted as being of more than one possible speech act (e.g. is it intended as a command or as advice?), or, if the utterance contains pronouns, demonstratives, or definite descriptions that do not uniquely identify their referents, it may be the case that which of the possible interpretations is correct can only be decided on the basis of recognition of the utterer's intention, and the speaker intends his utterance to be so disambiguated by his audience.

Bennett correctly realizes that not all occurrences of these kinds of ambiguity demand Gricean diagnoses of speaker's intentions. In cases where the context of utterance in conversation makes all but one interpretation nonsensical, or at least enormously unlikely, or where recognized intonation-contours or other characteristics of the utterance or speech-situation designate the intended speech act, or where one of our "abundance of non-Gricean ways" (LB, p.169) of sorting out intended referents come into play, a hearer may correctly interpret an utterance in virtue of general semantic properties of the utterance and its context and be intended to do so, or assumed to do so, by the speaker. (A speaker may not know his utterance is ambiguous in some way. This is irrelevant to the concerns at hand. If he doesn't know a problem exists, he can't have intentions as to how it should be solved.)

So, the cases that, in virtue of their ambiguity, are to show that a good portion of our language use involves

Gricean speaker's intentions are those in which there are no properties of the utterance that are conventionally used to disambiguate utterances of that kind, or properties of the context of utterance which would give reason to accept one interpretation as correct, and in which the speaker expects the hearer to recognize his intention to get him to believe a certain thing, and thereby come to have the appropriate belief. I can think of two ways in which this might occur. In one, the speaker may rely on a non-conventional way of disambiguating an utterance, and expect the hearer to infer his intention on the basis of that non-conventional signal. This case has few implications for the nature of the intentions and expectations in virtue of which we generally disambiguate utterances, insofar as we don't usually use non-conventional means of disambiguating utterances. Usually, it seems, we are pretty much unaware of ambiguity, trusting to context to sort things out. (If a non-conventional method of disambiguating utterances catches on, say because we happen upon some new ambiguity of speech act, and need to solve it, it becomes conventional, and won't do what Bennett wants it to.) The other situation I can think of in which a speaker expects his intention to get a hearer to recognize his intention to have the hearer form a certain belief, without there being properties of the utterance or its context which would suggest which interpretation was the correct one, is one in which the speaker has a completely groundless expectation. What is

said must be literally ambiguous, ambiguous in context, and not disambiguated by such things as facial expression or tone of voice, and the speaker must still expect his intended message to be recognized. Perhaps two persons are talking over a fence (they can't see each other's bodies) wearing welding masks (they can't see each others faces) where there is a great deal of noise (they are yelling as loud as they can in order to be heard) in an area where there are large heavy objects flying about (they often warn each other by yelling "Duck!"), looking into the sky looking for ducks. One yells "Duck!", expecting the other to recognize his intention to get the other to believe he should duck. It seems that any evidence that would give the hearer a clue as to what the speaker intended him to believe would also be evidence (on general semantic considerations) for a conclusion about a correct interpretation that did not depend on recognition of the speaker's intention. (We must be careful here: if the hearer of an ambiguous sentence forms a belief on the basis of general semantic considerations, he maybe does something we call "recognizing the speaker's intention"; although his recognizing the speaker's intention is not then part of his reason for holding the belief. His reasons are the properties of the utterance and its context through which he understood the utterance.) Bennett seems to require a case where the speaker expects his intention to be recognized, even while the hearer has no evidence as to what that intention might

be. While such speakers might be possible, I don't think they demonstrate much about our normal language use.

It seems that a large number of the ambiguous utterances we may have to deal with do not demand our doing any Gricean problem-solving. This is only to be expected, it seems to me. In any type of recurring situation in which we found ourselves a little baffled by what it was others were trying to get across, there would be pressure on us to develop methods of distinguishing possible messages on the basis of the ways and circumstances in which they were uttered. We have been here a long time, and it is no wonder most of the things we say are effectively unambiguous.

The second aspect of normal language use that Bennett claims demonstrates the existence of Gricean speaker's intentions in normal language use are injunctions that are intended by the speaker to be complied with in virtue of the hearer's altruism (LB, p.164ff). His case for claiming that expectations of altruistic compliance are Gricean involves a bit of subterfuge, however.

When an utterer U enjoins an audience A to do X, when there is no benefit for A in his doing X, it is often the case that U expects A's reason for doing X to be that, for A, satisfying certain wishes of persons or of U in particular is a good in itself. In such a case "A's basic reason for doing X is that U wants him to" (LB, p.165) that is, A does X in virtue of his recognition that U wants him to do, and that is what U intends. According to Bennett's

application of Grice's analysis of meaning to injunctions U has meant (Gricey) that A should do X. The subterfuge lies in Bennett's application of Grice's analysis of meaning to injunctions. Bennett follows Grice and says that "U meant that A is to do X if U acted with the intention

(i) that A should do X

(ii) that A should be aware of intention (a)

(iii) that the awareness mentioned in (a) should be part of A's reason for doing X." (LB, p.13)

Thus, one does not, as one might think one does, intend someone to understand what one is telling them to do, and expect them to do it on the basis of their recognition of what it is one wants them to do (when their altruism is supposed to provide their motivation). Rather, it is in virtue of their expected recognition of one's expectation one is supposed to mean anything at all, when enjoining someone to do something (when their altruism is expected to provide their motivation). Thus, the mechanisms by which one means what one does when one says "Close the door" are, by Bennett's account, supposedly entirely different when one says it to one's friend as opposed to when one says it to a subordinate, for example. I find this implausible.

The characteristically linguistic element of injunctions, whatever the expected reason for compliance can be given an account entirely in terms of their conventional meaning. With respect to injunctions in a language, we can say that someone is trusting in L, if they consider doing

what they are told to do or asked to do by a speaker of L, and truthful in L if they expect their hearers to do so. The reasons why they might do so, or be expected to do so may then be considered to be characteristics of the particular case, not essentially related to the conventional role of the utterance. In cases where the audience's altruism is expected to provide the motivation for compliance, the case will take on a Gricean character, but this character will not relate to the meaning of the utterance. The subterfuge in Bennett's account of injunction-meaning is the collapsing of the distinction between a communicative element and the intentions and beliefs involved in that, and the intentions and beliefs which come into play once someone knows what is wanted of them. The first has to do with the intentions involved in language use, the second does not. That the recognition of other's intentions plays a certain role in our social interactions does not mean that the recognition of other's intentions plays a role in the communication in virtue of which those interactions occur.

VII. Concluding Remarks

The individual is, in his future and in his past, a piece of fate, one law more, one necessity more for everything that is and everything that will be. To say to him 'change yourself' means to demand that everything should change, even in the past.

--Nietzsche, Twilight of the Idols

I'd like to describe what I take myself to have shown, and attempt to give some idea why someone would be concerned to have shown it. Two separable claims have been made. The bulk of this thesis has been intended to show that our conventions involve an essential element of the habitual acceptance of normality, and to apply this in refutation of two doctrines that have claimed that our conventional language use is not best understood as doing so. Also, it has been claimed that language is misrepresented if it is described in such a way as to imply an understanding on our part of a problem we collectively face and have found a way to solve. As David Lewis says, "Language is only one among many activities governed by conventions that we did not create by agreeing and cannot describe' (C, p.3).

Together, what do these imply? That our understanding is not required for the continuing operation of a convention, or for its success if it is a convention similar to language (one of the "others of this type"). Much of the problem of how to live together is solved, but not by us.

Both Grice and Bennett take positions that imply a realm of inference and understanding underlying the most basic processes via which we live with one another. I do not see that this is useful. I wonder, when we subtract from a man or a woman their part in our conventions, what can be thought to remain. Traditionally an appeal to convention has served to mark a person's ties to others and to the past. When applied to those things for which "convention" takes on a sense of metaphor, the use of the word has shifted the focus of a search to understand man away from what can be truly said of him independent of his being a social fact. I found it odd that an analysis of our concept of convention should serve the purposes of those who would deny this. I think I've shown that it should not.

The shameful secret of this thesis is its implicit conservatism. Without the man behind the scenes, the extent to which it is obvious that it is true that "We might change our conventions if we like" (C, p.1) decreases. If we take those things which permit our living together in virtue of our ability to conform, and not necessarily in virtue of our ability to understand, to be conventions, it is no longer clear that we would necessarily comprehend what a change in our conventions would consist in. In such a case, it may be closer to the truth to say we can cease to co-ordinate our actions if we like.

VIII. SELECT BIBLIOGRAPHY

- BENNETT, JONATHAN "The Meaning-Nominalist Strategy", Foundations of Language, Vol. 10, No. 1 (May 1973), pp. 141-168.
- BENNETT, JONATHAN Linguistic Behavior. Cambridge: Cambridge University Press, 1976.
- BIRO, JOHN "Intentionalism in the Theory of Meaning", The Monist, Vol. 62, No. 2 (April 1979), pp. 238-258.
- GILBERT, MARGARET "Convention and Game Theory", Synthese, Vol. 46, 1981.
- GRICE, H.P. "Meaning", The Philosophical Review, Vol. 66, No. 3 (1957), pp. 377-388.
- GRICE, H.P. "Utterer's Meaning, Sentence Meaning, and Word-Meaning", Foundations of Language, Vol. 4, No. 3 (August 1968), pp. 225-242.
- GRICE, H.P. "Utterer's Meaning and Intentions", The Philosophical Review, Vol. 78, No. 2 (April 1969), pp. 147-177.
- JAMIESON, D. "David Lewis on Convention", Canadian Journal of Philosophy, Vol. 5, No. 1 (September, 1975), pp. 73-81.
- LEWIS, DAVID K. Convention: A Philosophical Study. Cambridge, MA: Harvard University Press, 1969.
- LEWIS, DAVID K. "Languages and Language", Minnesota Studies in the Philosophy of Science, Vol. VII, ed. by Keith Gunderson. Minneapolis: University of Minnesota Press, 1975.
- LEWIS, DAVID K. "Reply to Jamieson", Canadian Journal of Philosophy, Vol. VI, No. 1 (March 1976), pp. 113-120.
- SCHELLING, THOMAS C., The Strategy of Conflict. Cambridge, MA: Harvard University Press, 1960.
- SCHIFFER, STEPHEN R. Meaning. Oxford: Oxford University Press, 1972.
- ULLMANN-MARGALIT, EDNA The Emergence of Norms. Oxford: Oxford University Press, 1977.

B30370